

PATENT APPLICATION

METHOD FOR ANALYZING PERFORMANCE INFORMATION

Inventors: **Kei TAKEDA**
Citizenship: Japan

Takato KUSAMA
Citizenship: Japan

Assignee: Hitachi, Ltd.
6, Kanda Surugadai 4-chome
Chiyoda-ku, Tokyo, Japan
Incorporation: Japan

Entity: Large

TOWNSEND AND TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(415) 576-0200

- 1 -

METHOD FOR ANALYZING PERFORMANCE INFORMATION

BACKGROUND OF THE INVENTION

The present invention relates to a system for displaying/analyzing performance information of servers and storage devices.

5 In recent years, configurations in which a plurality of servers share a single storage device have increased because of the development of the SAN (Storage Area Network) technique. The reason why such configurations are increasing is that the localization
10 of the operational business such as backup and fault control is facilitated and the operation management cost can be compressed.

On the other hand, however, since a plurality of servers share a single storage device, volume
15 assignment and storage performance management become complicated. A method for simplifying the volume assignment in the SAN environment is disclosed in US 2003/0005119A1.

In computer systems, the storage device is a
20 low speed storage medium as compared with caches and memories in the CPU. Often, therefore, the storage device becomes a bottleneck for the whole system. Therefore, it is important to acquire various performance index values in the construction phase and
25 the operation phase and make an analysis to determine

whether predetermined performance values have been obtained so that the storage device may exhibit as high performance as possible.

In a configuration in which a plurality of 5 servers share and use a single storage device, I/O transactions from a plurality of servers overlap on a specific resource within the storage device and consequently serious performance degradation is caused in some cases.

10 As an example, I/O loads from a plurality of servers concentrate on a specific physical disk group within a storage device, and the I/O performance seen from the server side is not sufficiently exhibited due to the I/O contention.

15 Owing to a tool provided by the storage device, the I/O quantity processed by a specific resource within the storage device can be perused/analyzed. In this method, however, it is difficult to discriminate the cause of an I/O 20 contention performance problem caused by setting on the outside of the storage device, i.e., specifically a server group that burdens an I/O load on a specific resource within the storage device or a portion that becomes a bottleneck of the performance.

25 SUMMARY OF THE INVENTION

An object of the present invention is to provide a method for solving the above-described

problem and facilitating the discrimination of a server group that is causing I/O contention on a resource in a storage device and a portion that has become a bottleneck in performance.

5 In accordance with the present invention, a performance information display method using a computer includes the steps, in the computer, of reading out information data of a storage device previously stored in a storage device and information data of a plurality 10 of devices utilizing the storage device, displaying an identifier of the storage device and identifiers of a plurality of devices utilizing the storage device on a screen on the basis of the information data read out, accepting a command to select the displayed identifier 15 of the storage device, and displaying performance information data of the devices utilizing the selected storage device in association on the basis of the accepted command and the information data read out.

 The present invention provides a method of 20 taking a specific resource within the storage device as a starting point and displaying at a stroke performance index values of a server group that burdens the load on the resource. Specifically, a performance management tool according to the present invention collects 25 mapping information between the storage device and the servers and performance information of the server group. If there is a command from the user, the performance management tool searches the mapping

information between the storage device and the servers for the server group utilizing a specific resource within the storage device, and creates a report of I/O contention narrowed down to the performance data of the 5 server group.

Other objects, features and advantages of the invention will become apparent from the following description of the embodiments of the invention taken in conjunction with the accompanying drawings.

10 BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram showing a configuration of a SAN system in an embodiment;

Fig. 2 is a diagram showing a volume provisioning function of a storage device and a volume 15 mount function of a server;

Fig. 3 is a diagram showing an internal structure of a performance information repository;

Fig. 4 is a diagram showing a cooperation method between a performance information collection 20 program and a performance information analysis program;

Fig. 5 is a diagram showing data items stored in a server performance information table;

Fig. 6 is a diagram showing data items stored in a server storage mapping table;

25 Fig. 7 is a diagram showing data items stored in a storage RAID group performance information table;

Fig. 8 is a diagram showing data items stored

in a storage port performance information table;

Fig. 9 is a diagram showing SCSI inquiry response information;

Fig. 10 is a flow chart showing processing of
5 a server performance information collection program;

Fig. 11 is a flow chart showing processing of
a storage performance information collection program;

Fig. 12 is a diagram showing an internal
structure of a performance information analysis program
10 and schematic views that are output by the performance
information analysis program;

Fig. 13 is a diagram showing an embodiment of
a storage performance information view;

Fig. 14 is a diagram showing an embodiment of
15 a related server performance information view;

Fig. 15 is a diagram showing an embodiment of
a related server performance history information view;

Fig. 16 is a flow chart showing processing of
a storage performance information display subprogram;

20 Fig. 17 is a flow chart showing processing of
a related server performance information display
subprogram; and

Fig. 18 is a flow chart showing processing of
a related server performance history information
25 display subprogram.

DETAILED DESCRIPTION OF THE EMBODIMENTS

Hereafter, an embodiment of the present

invention in a SAN system in which a plurality of servers share a storage device will be described.

Examples of the SAN system in the embodiment described hereafter are simplified as compared with the 5 ordinary SAN system because functions and details that are not necessary for description of the present embodiment are omitted. However, the application range of the present embodiment is not limited thereby.

Fig. 1 is a diagram showing a configuration 10 of a SAN system in the present embodiment. The present SAN system includes a storage device A 120, a server group including a server A 110A and a server B 110B, a SAN 140 for connecting the server group and the storage device to each other, a management server 130, and a 15 LAN 150 for connecting the server group and the storage device to the management server 130. Each of the servers may be a computer in which a program is executed, a terminal device using the function of the computer, or a different one.

20 Each of the server group and the storage device A transmits its performance information to the management server 130 via the LAN 150, and the management server 130 stores the performance information in its internal secondary storage device.

25 The server A 110A is connected to the SAN 140 via an HBA 114. The server A 110A executes a volume mount function 111A for mounting a volume provided from the storage device A 120, a business program 112A for

conducting business processing by using the mounted volume, and a server performance information collection program 101 for collecting performance information of a server volume.

5 In the same way as the server A, the server B 110B is connected to the SAN 140 via an HBA, and the server B 110B executes a volume mount function, a business program, and a server performance information collection program. Here, the SAN and LAN are used as 10 example of the networks. However, other networks such as NASs (Network Attached Storages) may also be used.

 The storage device A 120 is connected to the SAN 140 via a port A 123 and a port B 124. Within the storage device A 120, a volume provisioning function 15 121 for providing a volume for the SAN side and a storage performance information collection program 102 for collecting performance information of the storage device are executed.

 In the present embodiment, it is supposed for 20 brevity that one storage device is used, and the server group using the volume provided by the storage device A 120 includes only two servers: the server A 110A and the server B 110B. However, the present embodiment can also be applied to a configuration having more storage 25 devices and more servers.

 The management server 130 includes a secondary storage device 161, a main storage device 162, a communication device 163, a CPU 164, a display

device 165, a keyboard 166 and a mouse 167. Although not illustrated, the management server 130 also includes functions such as a file system required for a server to recognize data stored in the storage device.

5 A performance information repository 103 for preserving performance information of the server group and the storage device A 120 is stored in the secondary storage device 161. An operating system 168 is stored in the main storage device 162. Within the operating
10 system 168, a performance information analysis program 104 for analyzing performance information stored in the performance information repository 103 and displaying a result of the analysis on the display device 165.

Fig. 2 is a diagram showing the volume
15 provisioning function in the storage device A 120 and the volume mount function in the server group.

First, the volume provisioning function will now be described. The volume provisioning function 120 in the present embodiment provides a function of
20 enhancing the usability of physical disk groups in the storage device A 120 by using a configuration of the RAID (Redundant Array of Independent Disks), slicing a RAID group, which is virtual disks having a RAID configuration, in sizes that can be easily used by
25 higher rank servers, and making logical volumes obtained as a result of the slicing usable from the SAN side via a port. By thus combining a plurality of disks, it becomes possible to make the disk access fast

and enhance the fault resistance against failures of disks.

Specifically, the storage device A 120 includes eight physical disks, i.e., a physical disk A 5 240A, a physical disk B 240B, a physical disk C 240C a physical disk D 240D, a physical disk E 240E, a physical disk F 240F, a physical disk G 240G and a physical disk H 240H. In the present embodiment, the volume provisioning function 121 creates two RAID 10 groups, i.e., a RAID group A 230A and a RAID group B 230B by using these physical disks.

The RAID group is the unit obtained by regarding a plurality of physical disk groups having a RAID configuration as one virtual disk.

15 Furthermore, in the present embodiment, the volume provisioning function 121 logically slices the created RAID group, and creates six logical volumes as volumes that can be easily used by higher rank servers. In other words, the volume provisioning function 121 20 creates a logical volume A 220A, a logical volume B 220B and a logical volume C 220C from a RAID group A 230A, and creates a logical volume D 220D, a logical volume E 220E and a logical volume F 220F from a RAID group B 230B. Furthermore, the volume provisioning 25 function 121 opens the logical volumes to ports and makes the logical volumes accessible from the SAN 140 side.

For example, if an I/O request is issued to

the logical volume A by a server via the port A, the I/O request is temporarily converted to an I/O request for a corresponding slice portion within the RAID group A, and in addition it is further converted to an I/O request extending over the physical disk A 240A to the physical disk D 240D, which physically form the RAID group A, by the volume provisioning function 121, and processed.

The volume mount function will now be described. The volume mount function 111A in the server A 110A mounts volumes opened to the SAN 140 side from the storage device A 120 into the server A 110A, and makes the volumes usable from the business program 112A. The same holds true for the server B 110B as well. Typically, such a volume mount function is provided as a part of an operating system.

Finally, correlation between server side volumes and storage side volumes in the present embodiment will now be described. In the present embodiment, the server A 110a mounts the two volumes, i.e., the volume A 210A and the volume B 210B from the storage device A 120 and uses them. The volume A 210A and the volume B 210B correspond to the logical volume A 220A and the logical volume B 220B in the storage device A 120, respectively.

The server B 110B mounts a volume C 210C from the storage device A 120, and uses it. The volume C 210C corresponds to the logical volume C 220C in the

storage device A 120. All of the logical volumes A, B and C are logical volumes sliced from the RAID group A 230A. In the present embodiment, therefore, I/O loads of the server A and the server B become loads of the 5 same RAID group A, and it can be said that physical contention might occur between I/Os in the situation.

Fig. 3 is a diagram showing an internal structure of the performance information repository 103 stored in the secondary storage device 161 in the 10 management server 130. The performance information repository 103 includes four tables. A server performance information table 310 is a table storing performance information for the volumes in the server group. A server storage mapping table 320 is a table 15 storing mapping information among the volumes in the servers, the ports in the storage device, the logical volumes and the RAID groups. A storage RAID group performance information table 330 is a table storing performance information for RAID groups in the storage 20 device. A storage port performance information table 340 is a table storing performance information for ports in the storage device.

Fig. 4 is a diagram showing a cooperation method between the performance information collection 25 program 101 and the performance information analysis program 104.

As described earlier, the server performance information collection program 101 is operating on the

server A 110A and the server B 110B. The server performance information collection program 101 collects performance information for the volumes on the server, and adds records to the server performance information 5 table 310 and the server storage mapping table 320 on the performance information repository 103.

The storage performance information collection program 102 is operating in the storage device A 120. The storage performance information 10 collection program 102 collects performance information for the ports and the RAID groups in the storage device, and adds records to the storage RAID group performance information table 330 and the storage port performance information table 340 on the performance 15 information repository 103.

The performance information analysis program 104 is operating on the management server 130. In response to a request from the user, the performance information analysis program 104 acquires and analyzes 20 the performance information data on the performance information repository 103, and displays its result on the display device 165.

Fig. 5 is a diagram showing data items stored in the server performance information table 310 in 25 detail. The present table is a table storing I/O performance information for the volumes on the servers, and it is a table to which records are added by the server performance information collection program 101.

In the present embodiment, the present table includes a server column 511 and a volume column 512, serving as columns for uniquely identifying servers and volumes.

The present table includes an IOPS (IO per second) column 513, a Read IOPS (IO per second) column 514 and a Write IOPS (IO per second) column 515 respectively representing the number of IO times, the number of read IO times and the number of write IO times per second for respective volumes.

The present table includes an Xfer (Transfer) column 516, a Read Xfer (Transfer) column 517 and a Write Xfer (Transfer) column 518 respectively representing the transfer data quantity, read transfer data quantity and the write data quantity per second for respective volumes. Finally, the present table 310 includes a TS column 519 representing time when the performance information was observed.

Fig. 6 is a diagram showing data items stored in the server storage mapping table 320 in detail. The present table is a table representing each volume in each server corresponds to which logical volume and which RAID group via which port in the storage device, and it is a table to which records are added by the server performance information collection program 101.

In the present embodiment, the present table includes a server column 611 and a volume column 612, serving as columns for uniquely identifying servers and volumes. The present table includes a storage column

613, a port column 614, a logical volume column 615, and a RAID group column 616, which represent each volume is using which one of storage devices, ports, and logical volumes and RAID groups.

5 In the present embodiment, it is supposed that the server storage mapping table 320 has only the latest mapping information. However, it is also possible to add a mapping information acquisition time column to the present table to provide the present
10 table with past history information of the mapping information, and thereby create a report having related server groups narrowed down on the basis of the past mapping information.

Fig. 7 is a diagram showing data items stored
15 in the storage RAID group performance information table 330 in detail. The present table is a table storing I/O performance information for respective RAID groups in respective storage devices, and it is a table to which records are added by the storage performance
20 information collection program 102. In the present embodiment, the present table includes a storage column 711 and an RAID group column 712, serving as columns for uniquely identifying RAID groups on storage devices. The present table includes an IOPS (IO per
25 second) column 713, a Read IOPS (IO per second) column 714 and a Write IOPS (IO per second) column 715 respectively representing the number of IO times, the number of read IO times and the number of write IO

times per second for respective RAID groups.

The present table includes an Xfer (Transfer) column 716, a Read Xfer (Transfer) column 717 and a Write Xfer (Transfer) column 718 respectively 5 representing the transfer data quantity, the read transfer data quantity and the write data quantity per second for respective RAID groups. Finally, the present table includes a TS column 719 representing time when the performance information was observed.

10 Fig. 8 is a diagram showing data items stored in the storage port performance information table 340 in detail. The present table is a table storing I/O performance information for respective ports in respective storage devices, and it is a table to which 15 records are added by the storage performance information collection program 102.

In the present embodiment, the present table includes a storage column 811 and a port column 812, serving as columns for uniquely identifying ports on 20 storage devices. The present table includes an IOPS (IO per second) column 813, a Max IOPS (IO per second) column 814 and a Min IOPS (IO per second) column 815 respectively representing the number of IO times, the maximum number of IO times and the minimum number of IO 25 times per second for respective ports. The present table includes an Xfer (Transfer) column 816, a Read Xfer (Transfer) column 817 and a Write Xfer (Transfer) column 818 respectively representing the transfer data

quantity, the maximum transfer data quantity and the minimum data quantity per second for respective ports. Finally, the present table includes a TS column 819 representing time when the performance information was 5 observed.

Fig. 9 is a diagram showing SCSI inquiry response information. SCSI inquiries can be issued from the server to respective volumes. The issued SCSI inquiry is transferred to a storage device via the SAN. 10 The storage device returns SCSI inquiry response information shown in Fig. 9 to the server side as its response.

In the present embodiment, the server performance information collection program 101 issues 15 SCSI inquiries to volumes on respective servers, and acquires information of a storage device side of connection destination by utilizing returned SCSI inquiry response information. Furthermore, in the present embodiment, the SCSI inquiry response 20 information holds a storage device name, a port name, a logical volume name and a RAID group name of the connection destination in a connection destination storage column 911, a connection destination port column 912, a connection destination logical volume 25 column 913 and a connection destination RAID group column 914.

Here, an example in which various kinds of information are collected by using an SCSI command is

shown as an example. However, information may be collected by using a method other than this.

Fig. 10 is a flow chart showing processing of the server performance information collection program

5 101. The server performance information collection program 101 is a program for acquiring performance information and storage device information of connection destination for all volumes on a server on which the present program operates and adding records
10 10 to the performance information repository 103.

At step 1001, initialization processing for repetitively conducting processing for all volumes in the server is conducted.

At step 1002, the OS is requested to acquire
15 the performance information of the volume and its result is acquired. In the present embodiment, the number of IO times, the number of read IO times and the number of write IO times per second, the transfer data quantity, the read transfer data quantity and the write
20 data quantity per second are acquired for respective volumes, as an example.

At step 1003, the performance information of the volume acquired at the step 1002 is added to the server performance information table 310 together with
25 a time stamp.

At step 1004, a SCSI inquiry is issued for the volume and SCSI inquiry information 910 is acquired. Furthermore, information of the storage

device, port, logical volume and RAID group that are being utilized by the volume is extracted from the acquired SCSI inquiry information 910.

At step 1005, connection destination 5 information of the volume acquired at the step 1004 is assed to the server storage mapping table 320. If mapping information of the volume already exists, it is deleted and then addition is conducted.

At step 1006, it is checked whether the 10 volume is the last volume on the server. If the volume is the last volume, the processing is finished. If there are still remaining volumes, the processing returns to the step 1002.

Fig. 11 is a flow chart showing processing of 15 the storage performance information collection program 101. The storage performance information collection program 101 is a program for acquiring performance information for all ports and RAID groups on the storage device on which the present program operates, 20 and adding records to the performance information repository 103.

At step 1101, initialization processing for repetitively conducting processing for all ports in the storage device is conducted.

25 At step 1102, the storage device is requested to acquire the performance information of the port and its result is acquired. In the present embodiment, the number of IO times, the maximum number of IO times and

the minimum number of IO times per second, the transfer data quantity, the maximum transfer data quantity and the minimum data quantity per second are acquired for respective ports, as an example.

5 At step 1103, the performance information of the port acquired at the step 1102 is added to the storage port performance information table 340 together with a time stamp.

10 At step 1104, it is checked whether the port is the last port on the storage device. If the port is the last port, the processing proceeds to step 1105. If there is a remaining port, the processing returns to the step 1102.

15 At step 1105, initialization processing for repetitively conducting processing for all RAID groups in the storage device is conducted.

20 At step 1106, the storage device is requested to acquire the performance information of the RAID group and its result is acquired. In the present embodiment, the number of IO times; the number of read IO times and the number of write IO times per second, the transfer data quantity, the read transfer data quantity and the write data quantity per second are acquired for respective RAID groups, as an example.

25 At step 1107, the performance information of the RAID group acquired at the step 1106 is added to the storage RAID group performance information table 330 together with a time stamp.

At step 1108, it is checked whether the RAID group is the last RAID group on the storage device. If the RAID group is the last RAID group, the processing is finished. If there is a remaining RAID group, the 5 processing returns to the step 1106.

Fig. 12 is a diagram showing an internal structure of the performance information analysis program 104 and schematically showing a view that is output by the present program.

10 The performance information analysis program 104 includes three internal subprograms, i.e., a storage performance information display subprogram 1201, a related server performance information display subprogram 1202 and a related server performance 15 history information display subprogram 1203.

The storage performance information display subprogram 1201 is a program for providing the user with performance information of ports and RAID groups in the storage device by displaying a storage 20 performance information view 1210.

The related server performance information display subprogram 1202 is a program for narrowing down volumes to volumes on the server side utilizing ports or RAID group resources of a certain specific storage 25 side and displaying performance information of a plurality of server side volumes by displaying a related server performance information view 1220.

The related server performance information

view 1220 is started from the storage performance information view 1210 according to a user's request. By watching the related server performance information view 1220, it becomes easy for the user to discriminate 5 a server that burdens a load on a specific resource in the storage device. And by perusing the performance information of the server group at a stroke, a server that especially burdens load on the resource can be discriminated.

10 The related server performance history information display subprogram 1203 is a program for displaying performance history information of a server group using a specific resource in the storage device in a time series form in the order of time when 15 performance information was observed by displaying a related server performance history information view 1230. The related server performance history information view 1230 is started from the related server performance information view 1220 according to a 20 user's request.

Fig. 13 is a diagram showing an embodiment of the storage performance information view 1210. The present view includes a storage object selection region 1310, a performance display date and hour input region 1320, a selected object performance information display region 1330 and a related server performance display button 1340.

In the present embodiment, a storage device

and resources within the storage device are displayed in a tree form in the storage performance information view 1210. Specifically, in Fig. 13, a storage device A node 1311 is displayed as a parent node of the tree, 5 and RAID groups and ports are displayed as child nodes like a RAID group A node 1312, a RAID group B node 1313, a port A node 1314 and a port B node 1315. The user can select one of the child nodes in the present storage object selection region, and performance 10 information concerning the selected object is displayed in the selected object performance information display region 1330.

In the present embodiment, kinds of resources within the storage to be subjected to the performance 15 analysis are two kinds, i.e., the RAID group and port. However, the present embodiment can also be applied to other resources such as cache memories in the storage device, and the application is not limited.

The performance display date and hour input 20 region 1320 is a region for the user to input the time when the performance information to be displayed. Performance information at the date and hour that has been input to the region is retrieved and displayed in the selected object performance information display 25 region 1330. As for the specification of the date and hour for displaying the performance information, the program may previously specify that the current time should be displayed, or performance information in a

predetermined interval may be displayed in time series order according to data accepted by the user. Other specification is also possible.

The selected object performance information
5 display region 1330 is a region for displaying performance information concerning an object selected in the storage object selection region 1310 as regards the date and hour that has been input in the performance display date and hour input region 1320.

10 In Fig. 13, it is supposed that the RAID group A node 1312 has been selected in the storage object selection region 1310 and "10 a.m., July 19, 2003" has been input as the performance display date and hour. Therefore, performance information of the
15 RAID group A at 10 a.m., July 19, 2003 is displayed in the selected object performance information display region 1330. Specifically, the selected resource name, i.e., "RAID group A" is displayed in a selected object display region 1337, and the performance information of
20 the RAID group A at 10 a.m., July 19, 2003 is displayed in an IOPS (IO per second) column 1331, a Read IOPS (IO per second) column 1332, a Write IOPS (IO per second) column 1333, an Xfer (Transfer) column 1334, a Read Xfer (Transfer) column 1335 and a Write Xfer (Transfer)
25 column 1336.

The related server performance display button 1340 is a button for requesting the display of the related server performance information view 1220. By

pressing the present button, it is possible to display the related server performance information view 1220 for displaying the performance information of a volume on the server side utilizing an object selected in the 5 storage object selection region 1310. On the inside, the present button starts the related server performance information display subprogram 1202.

Fig. 14 is a diagram showing an embodiment of the related server performance information view 1220.

10 The present view is a view that displays a list of servers and volumes utilizing a specific RAID group or port, and their performance information. By watching the present view, the user can discriminate a server and a volume that burden the load on a specific RAID 15 group or port.

The related server performance information view 1220 includes a related server performance information display region 1410 and a history information display button 1420.

20 The related server performance information display region 1410 is a region for displaying a list of related servers and volumes and their performance information. In the present embodiment, the related server performance information display region 1410 includes a server column 1421 and a volume column 1422 for specifying a server and a volume that utilize the RAID group or port. Furthermore, in the present embodiment, an IOPS (IO per second) column 1423, a

Write IOPS (IO per second) column 1424, a Read IOPS (IO per second) column 1425, an Xfer (Transfer) column 1426, a Read Xfer (Transfer) column 1427 and a Write Xfer (Transfer) column 1428 are displayed as the 5 performance information of the volume in the related server performance information display region 1410.

The history information display button 1420 is a button for requesting the display of the related server performance history information view 1230. By 10 pressing the present button, the related server performance history information view 1230 for historically displaying the performance information of the related server group can be displayed. On the inside, the present button starts the related server 15 performance history information display subprogram 1203.

Fig. 15 is a diagram showing an embodiment of the related server performance history information view 1230. The present view historically displays the 20 performance information of the related server group. The present view includes a related server performance history information display region 1510. Historical performance information of the related server is displayed in the related server performance history 25 information display region 1510 by broken line graphs as shown in Fig. 15.

In Fig. 15, the abscissa of the graph indicates time, and the ordinate of the graph indicates

performance, as an example. Performance history data for the volume A in the server A and performance history data for the volume C in the server B are displayed by superposition of broken line graphs, and a 5 server name and a volume name are added to each of the broken line graphs. However, this is an example, and performance information of devices other than servers may also be displayed.

For example, performance information may be 10 displayed by using one view every logical volume. In the case where volumes recognized by respective servers utilize the same RAID group, the performance information may be displayed by using a window every RAID group. Or other display methods may also be used.

15 It is also possible to partition one window into a plurality of parts and display various devices in a tree structure on a left-hand part of the window as represented by 1310 in Fig. 13. When the user has selected one of components in the tree structure shown 20 in 1310, performance history information display of related devices (for example, Fig. 15) may be conducted on a right-hand part of the window on the basis of the user's selection.

Information data of devices sharing a storage 25 device specified by the user may be superposed and displayed by using graphs having the same coordinate axes. The performance information of the specified storage device and performance information data of

three devices having the highest loads among devices utilizing the specified storage device may be displayed so as to be associated with each other. A different display method may also be used. When displaying a 5 plurality of performance information data, the performance information data may be displayed in one window so as to be superposed in the order of increasing value, the performance information data may be displayed in one window in accordance with a 10 predetermined order, or a plurality of windows may be displayed side by side in the horizontal direction or the vertical direction in association.

Thus, performance data of the devices and computers are not shown singly, but they are displayed 15 in association with performance information of related devices on a path between a server and a storage device and performance information of devices sharing the storage device. As a result, it becomes easy to discriminate a part that becomes a performance 20 bottleneck even in a complicated system environment.

It becomes possible to narrow down and display performance information data of related devices from among a large number of performance information data. It becomes easy for a human being to understand 25 the performance information.

Fig. 16 is a flow chart showing processing of the storage performance information display subprogram 1201. The present program is a program for displaying

the storage performance information view 1210.

The present program is broadly divided into two kinds of processing. Steps 1601 to 1603 are processing for drawing the storage object selection 5 region 1310. Processing ranging from the step 1604 to end processing is processing for drawing the selected object performance information display region 1330.

At step 1601, the storage RAID group performance information table 330 is scanned, and a 10 list of pairs of all RAID groups and storage devices to which the RAID groups belong is acquired.

At step 1602, the storage port performance information table 340 is scanned, and a list of pairs of all ports and storage devices to which the ports 15 belong is acquired.

At step 1603, all storage devices, RAID groups and ports are drawn in the storage object selection region 1310 on the basis of the list of RAID groups paired with storage devices and the list of 20 ports paired with storage devices acquired respectively at the steps 1601 and 1602.

At step 1604, date and hour information that has been input in the performance display date and hour input region 1320 is acquired.

25 At step 1605, a name of an object selected in the storage object selection region 1310 is acquired. For example, specification of an identifier of an RAID group, a port name, and a name and an identifier of a

storage related device is accepted.

At step 1607, on the basis of the name of the storage object acquired at the step 1605, it is determined whether the object is an RAID group or a 5 port. If the object is a RAID group, the processing proceeds to a step 1608. If the object is a port, the processing proceeds to step 1610.

At step 1608, the storage RAID group performance information table 330 is searched by using 10 the object name acquired at the step 1605 and the date and hour acquired at the step 1604 as a key, and performance information of the selected RAID group at the specified date and hour is acquired.

At step 1609, performance information of the 15 RAID group acquired at the step 1608 is drawn in the selected object performance information display region 1330.

At step 1610, the storage port performance information table 340 is searched by using the object 20 name acquired at the step 1605 and the date and hour acquired at the step 1604 as a key, and performance information of the selected port at the specified date and hour is acquired.

At step 1611, performance information of the 25 port acquired at the step 1610 is drawn in the selected object performance information display region 1330.

Fig. 17 is a flow chart showing processing of the related server performance information display

subprogram 1202. The present program is a program that is executed to draw the related server performance information view 1220 when the user has pressed the related server performance display button 1340 on the 5 storage performance information view 1210.

At step 1701, date and hour information that has been input in the performance display date and hour input region 1320 is acquired.

At step 1702, a name of an object selected in 10 the storage object selection region 1310 is acquired.

At step 1703, on the basis of the name of the storage object acquired at the step 1702, it is determined whether the object is an RAID group or a port. If the object is a RAID group, the processing 15 proceeds to a step 1704. If the object is a port, the processing proceeds to step 1705.

At step 1704, the server storage mapping table 320 is searched by using the object name acquired at the step 1702 as a key for the RAID group column, 20 and a list of volumes of servers utilizing the RAID group is acquired.

At step 1705, the server storage mapping table 320 is searched by using the object name acquired at the step 1702 as a key for the port column, and a 25 list of volumes of servers utilizing the port is acquired.

At step 1706, the server performance information table 310 is searched by using the date and

hour acquired at the step 1701 and the volume list of servers acquired at the step 1704 or 1705 as a key, and performance information of the volume group of the server side utilizing the RAID group or the port is 5 acquired.

At step 1707, performance information of the related server acquired at the step 1706 is drawn in the related server performance information display region 1410.

10 Fig. 18 is a flow chart showing processing of the related server performance history information display subprogram 1203. The present program is a program that is executed to draw the related server performance history information view 1230 when the user 15 has pressed the history information display button 1420 on the related server performance information view 1220.

At step 1801, a list of volumes of the server side displayed in the related server performance 20 information region 1410 is acquired.

At step 1802, the server performance information table 310 is searched by using the volume list of the server side acquired at the step 1801 as a key, and historical performance data of the volume 25 group is acquired.

At step 1803, the related server performance history information view 1230 is created, and historical performance data of related volumes acquired

at the step 1802 are drawn in the related server performance history information display region 1510.

The above-described processing and display transition are an example, and different view 5 transition may be conducted. For example, in the case where the user has specified the RAID group A node 1312 in the view shown in Fig. 13 with a mouse pointer or the like, load information of a plurality of servers utilizing the RAID group A may be displayed in a form 10 of superposed graphs (for example, as shown in Fig. 15). Furthermore, a port may be specified and load information of a server utilizing the port may be displayed on the basis of the specified port information. Or it is also possible to associate route 15 information utilizing the port with a storage device on the route (for example, in a form such as "server - volume name - switch (I/O port) - logical volume - RAID group - physical disk"), and display information relating to each of the storage devices. Different 20 display may also be conducted.

Furthermore, in the case where the fact that the I/O performance of a certain server has degraded is detected, it is also possible to regard the server degraded in performance as a reference point and 25 display performance information of ports utilized by the server and information of servers utilizing the ports. Or it is also possible to check logical volumes and RAID groups utilized by the server, and display

performance information of other servers sharing and utilizing the logical volumes and RAID groups.

As described above, according to the storage resource contention situation detection method, a specific resource within the storage device is regarded as the starting point, and performance index values of the server group (or various devices) that burdens the load on the resource are displayed at a stroke. In a configuration in which a plurality of servers share and utilize a single storage device, therefore, it is possible to easily detect concentration of I/O loads on a resource within the storage device and easily discriminate servers that burden the load.

Furthermore, in the case where a plurality of storage devices or a plurality of servers share various devices, it becomes easy to grasp the state of a device that becomes a bottleneck in performance of the shared devices by collecting and analyzing performance information of related devices.

According to the present invention, performance values of devices that burden the load on a storage device are displayed in association. As a result, it is possible to easily discriminate a server group that causes I/O contention on a resource in the storage device, or a part that becomes a bottleneck in performance.

It should be further understood by those skilled in the art that although the foregoing

description has been made on embodiments of the invention, the invention is not limited thereto and various changes and modifications may be made without departing from the spirit of the invention and the 5 scope of the appended claims.